

Maya: Next Generation Modeling & Simulation Tools for Global Networks

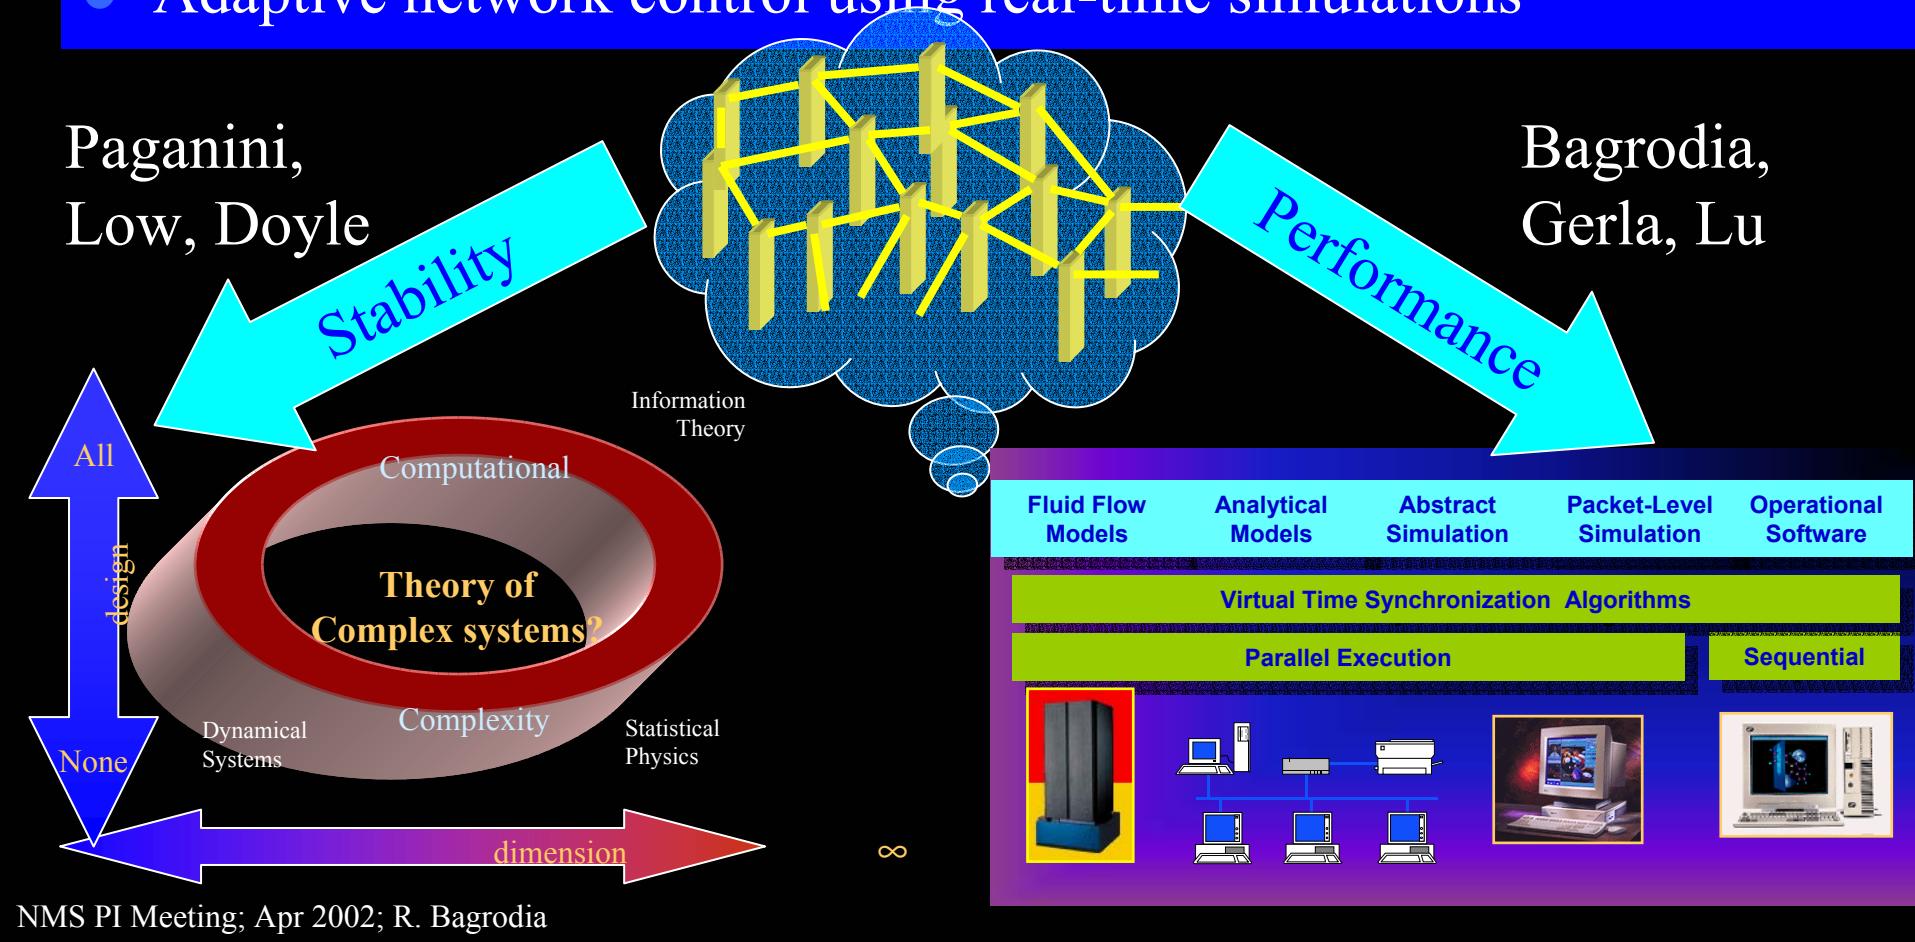
UCLA: R. Bagrodia, M. Gerla, S. Lu, F. Paganini,
M. Sanadidi, M. Takai

Caltech: J. Doyle, S. Low

DARPA PI Meeting, April 2002
Baltimore

Impact

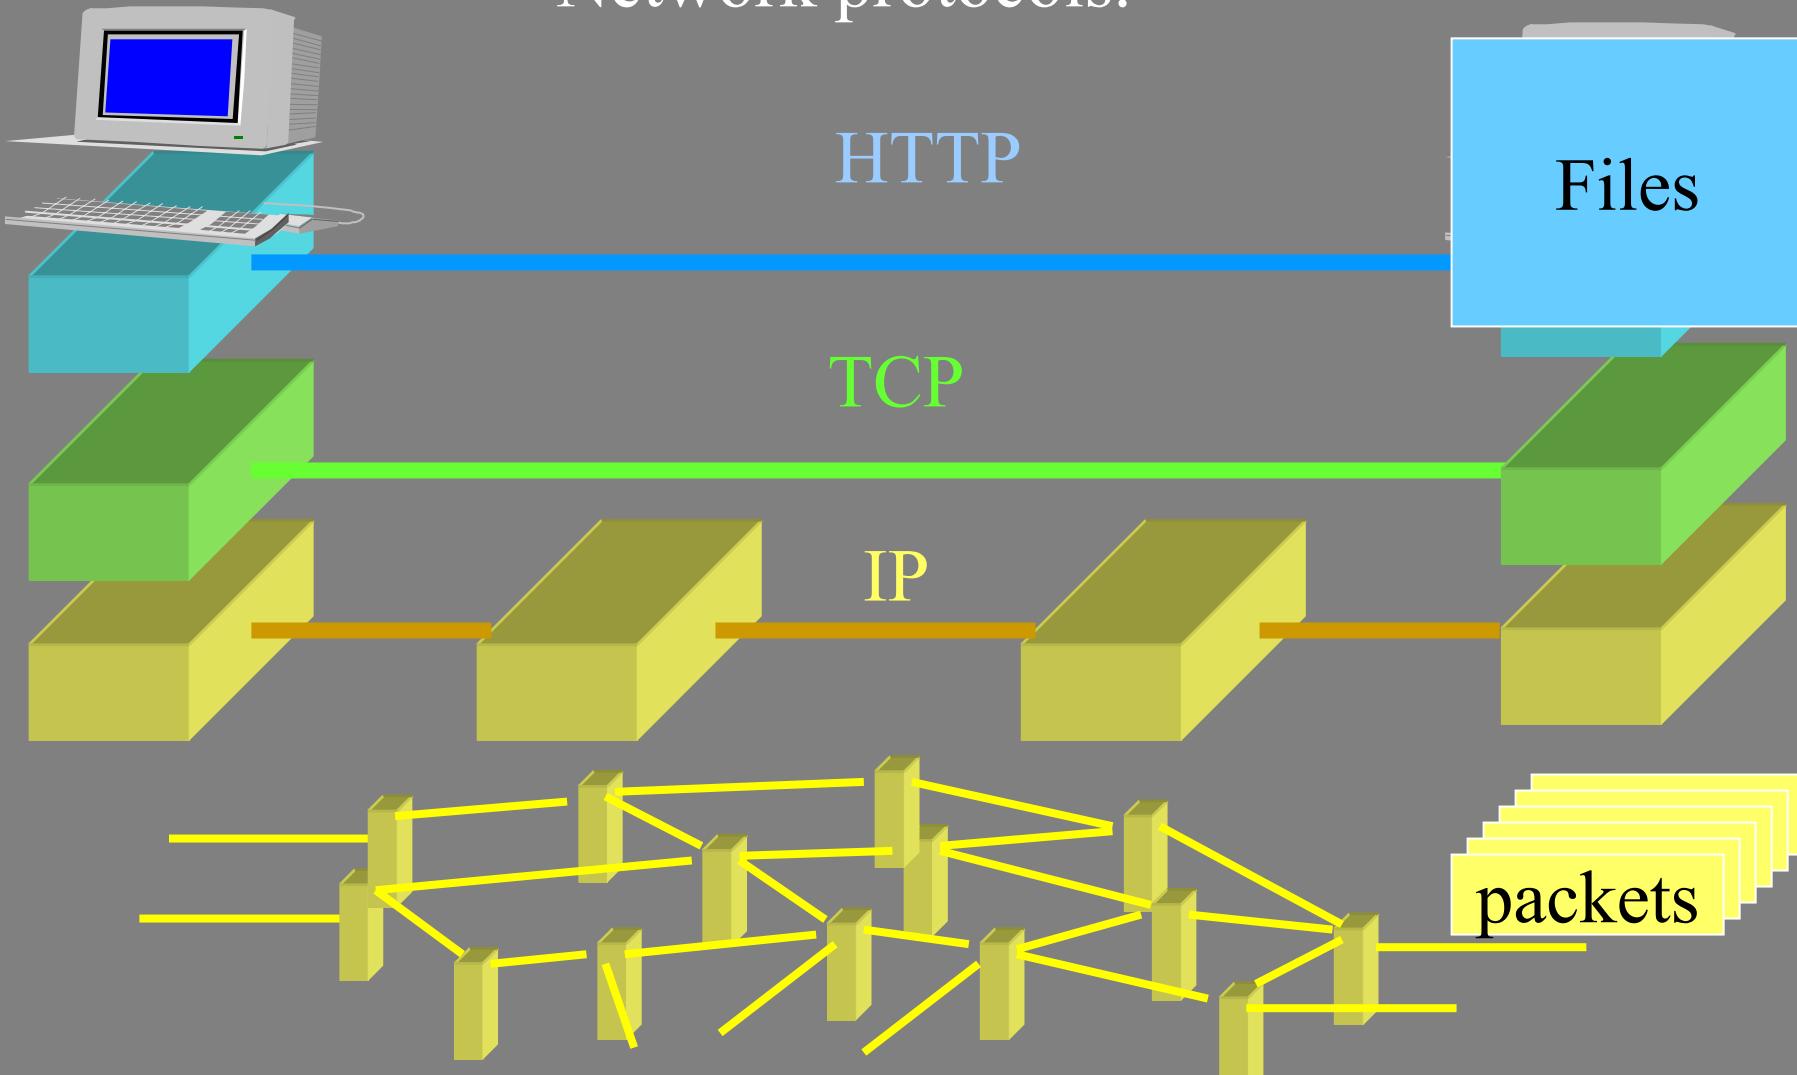
- A *multi-paradigm* integrated framework for network performance analysis: Maya
- A novel theory of network control and stability: HOT and FAST
- Adaptive network control using real-time simulations



Outline

- HOT traffic
- Duality model of TCP/AQM
 - Current protocols (Reno, Vegas, RED, REM/PI, AVQ...)
- FAST control
 - New protocols
- Theory of complex systems

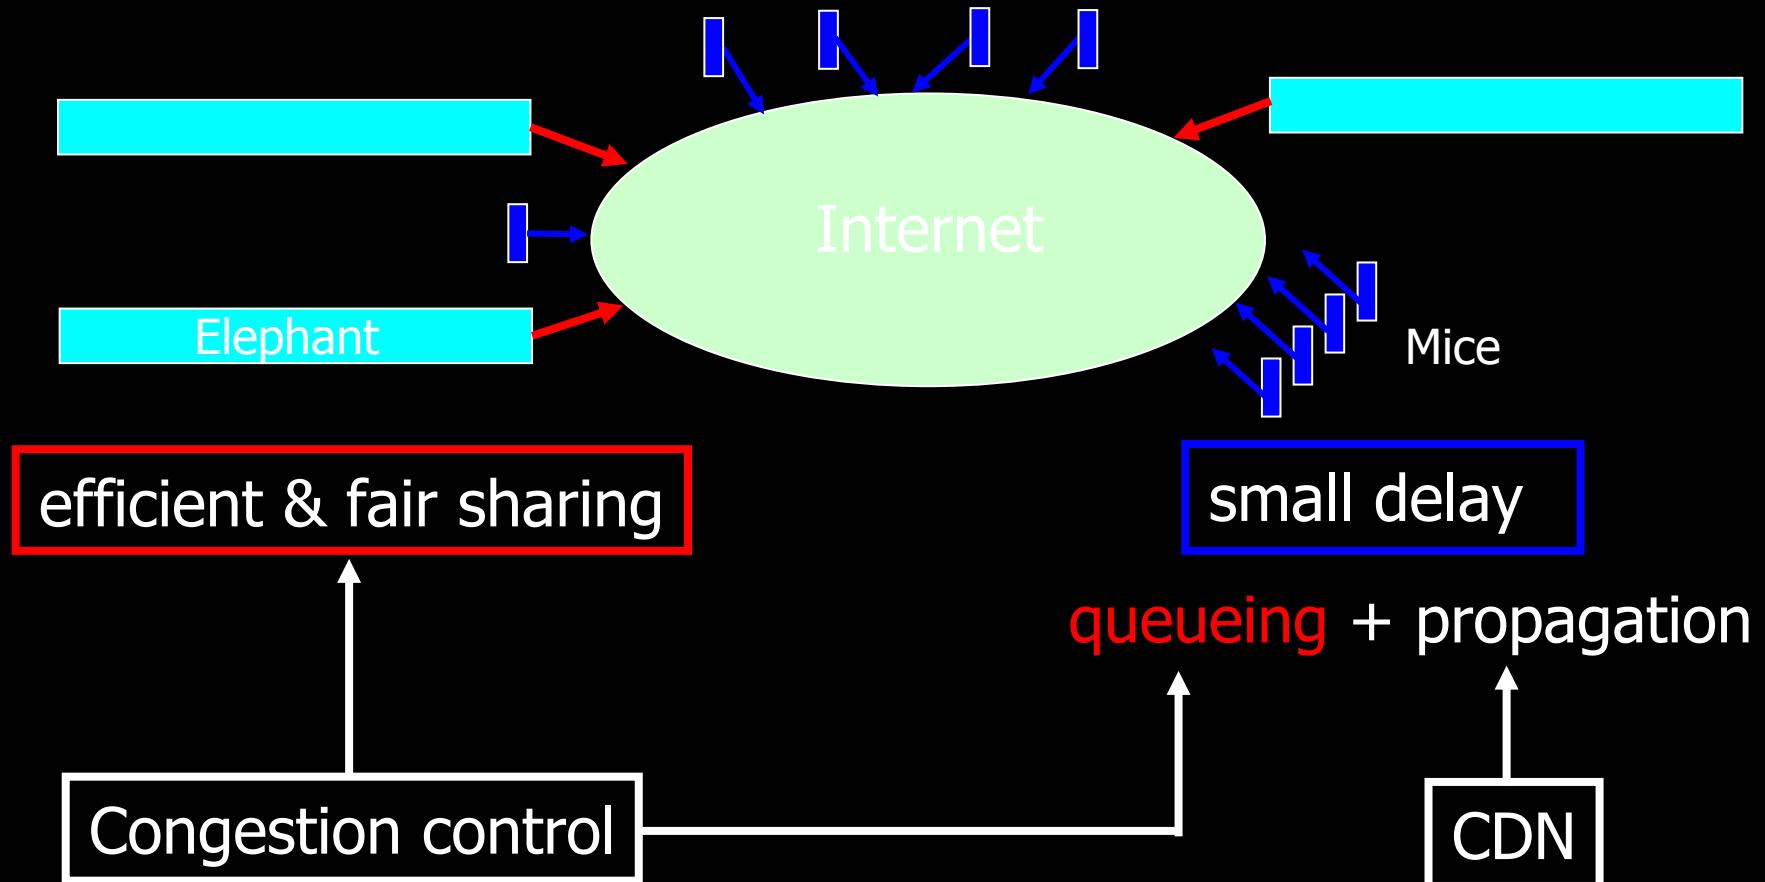
Network protocols.



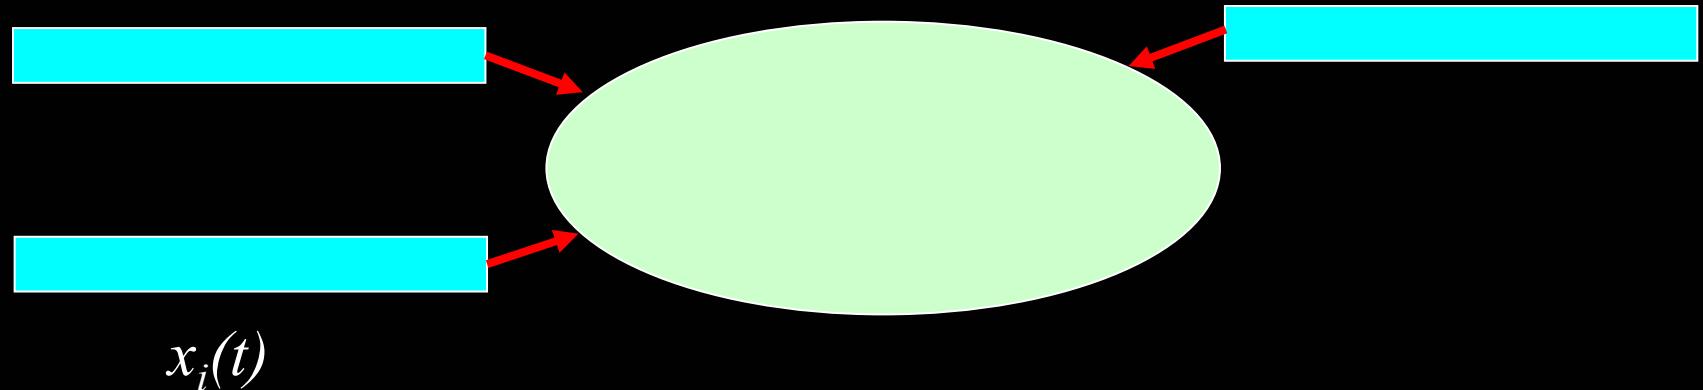
Routing
Provisioning

Congestion Control

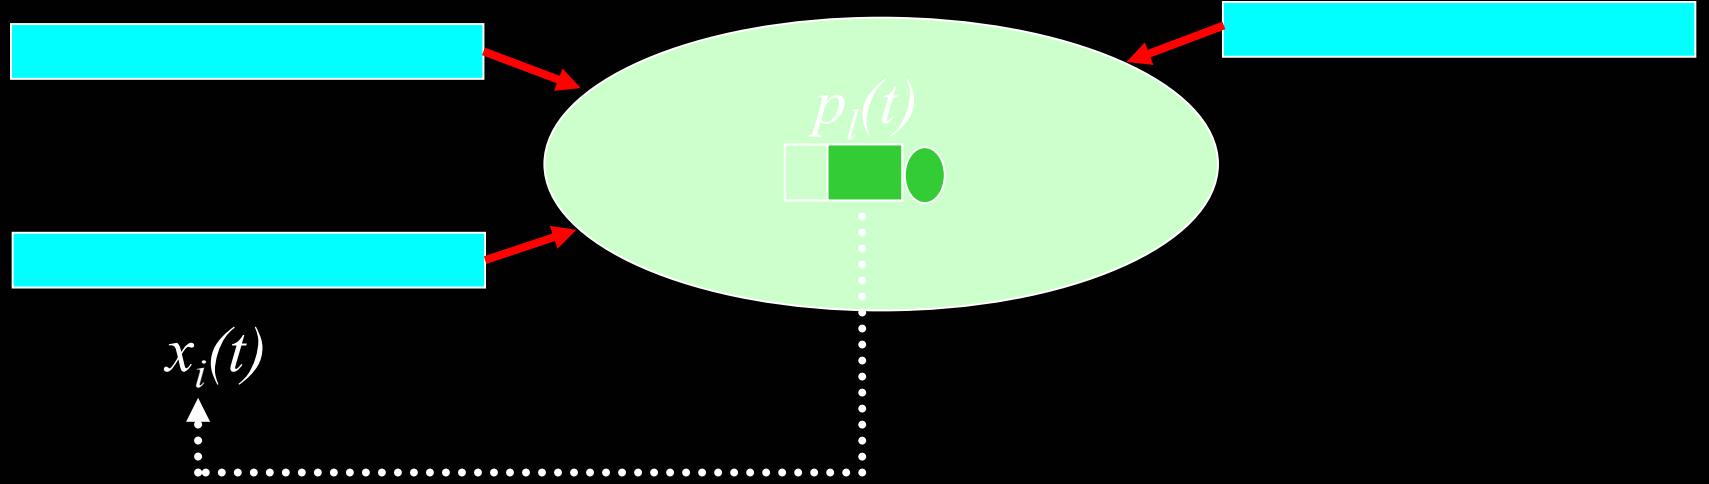
Heavy tail → Mice-elephants



TCP



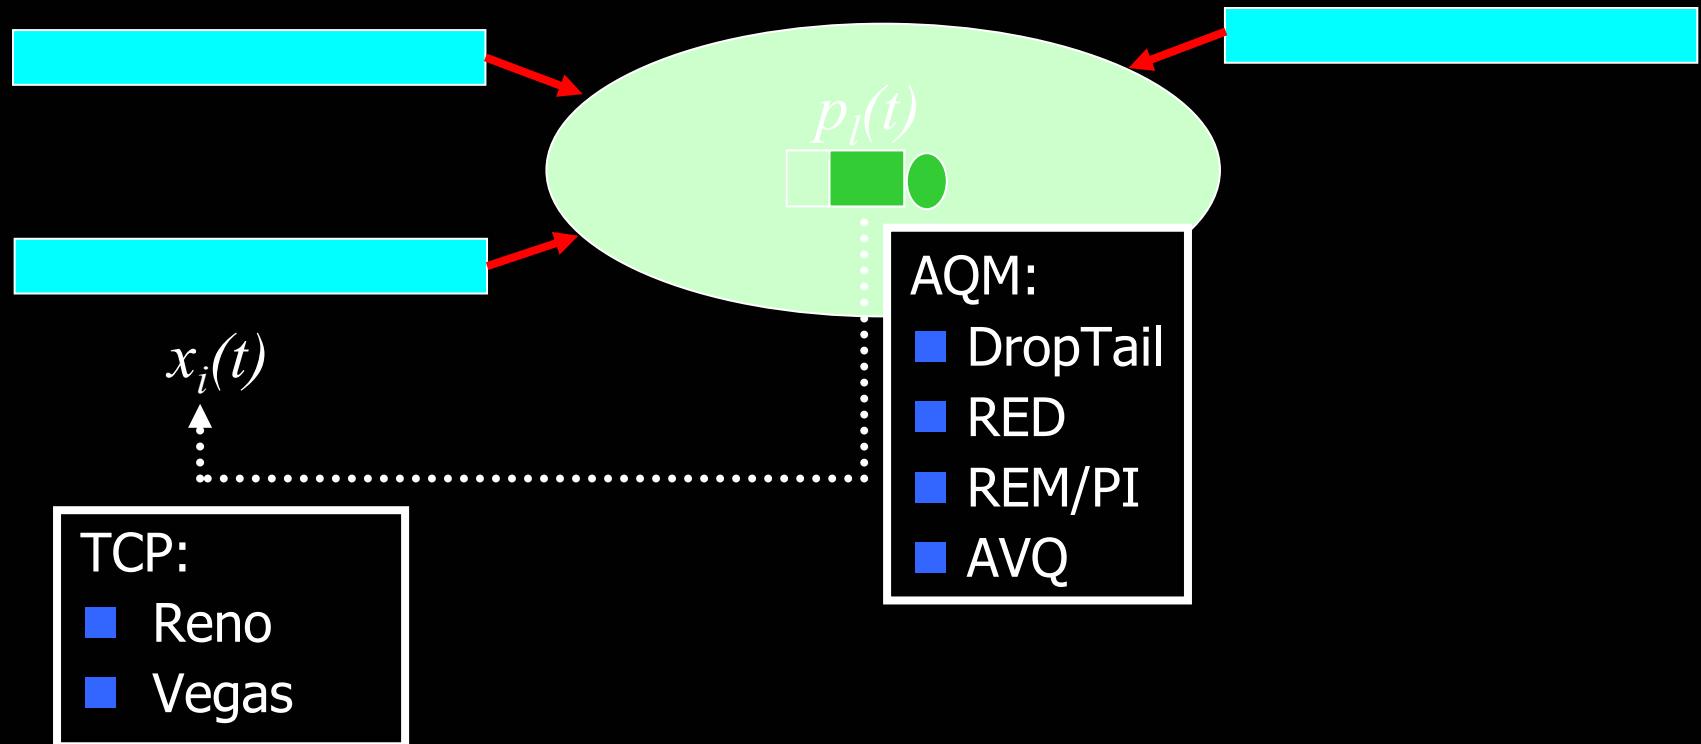
TCP & AQM



Example congestion measure $p_l(t)$

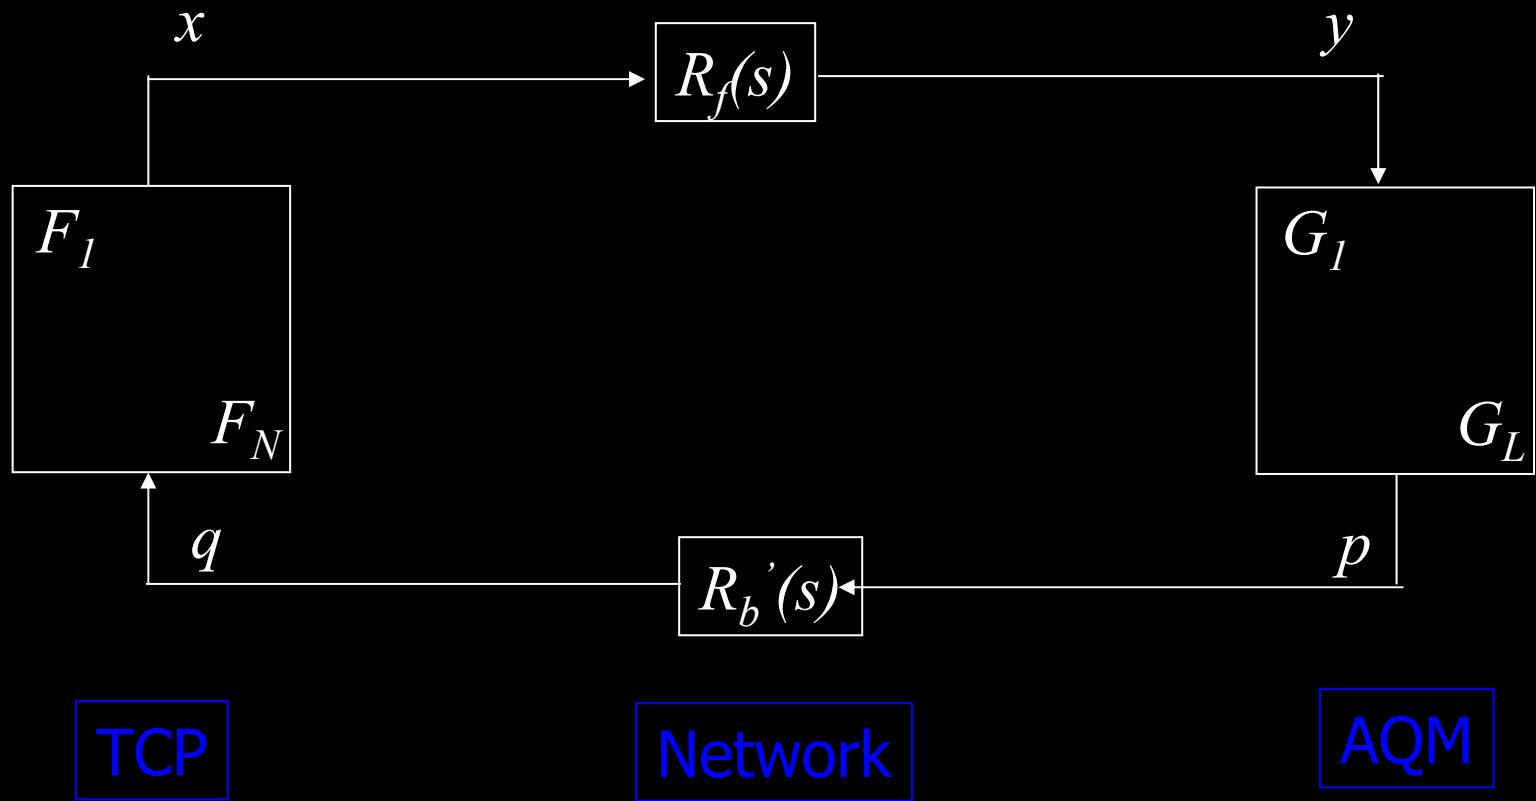
- Loss (Reno)
- Queueing delay (Vegas)

TCP & AQM



Model structure

Multi-link multi-source network



Overview

Protocol (Reno, Vegas, RED, REM/PI...)

$$x(t+1) = F(p(t), x(t))$$

$$p(t+1) = G(p(t), x(t))$$

Equilibrium

- Performance
 - Throughput, loss, delay
- Fairness
- Utility

Dynamics

- Local stability
- Cost of stabilization

Congestion Control

$$\text{Primal: } \max_{x_s \geq 0} \sum_s U_s(x_s) \quad \text{subject to } x^l \leq c_l, \quad \forall l \in L$$

$$\text{Dual: } \min_{p \geq 0} \quad D(p) = \left(\max_{x_s \geq 0} \sum_s U_s(x_s) + \sum_l p_l (c_l - x^l) \right)$$

Primal-dual algorithm:

$$x(t+1) = F(p(t), x(t))$$

$$p(t+1) = G(p(t), x(t))$$

Duality Model of TCP/AQM

- Source algorithm iterates on rates
- Link algorithm iterates on prices
- With different utility functions

Primal-dual algorithm:

$$x(t+1) = F(p(t), x(t))$$

Reno, Vegas

$$p(t+1) = G(p(t), x(t))$$

DropTail, RED, REM

Duality Model of TCP/AQM

(x^*, p^*) primal-dual optimal if and only if

$$y_l^* \leq c_l \quad \text{with equality if } p_l^* > 0$$

Primal-dual algorithm:

$$x(t+1) = F(p(t), x(t))$$

Reno, Vegas

$$p(t+1) = G(p(t), x(t))$$

DropTail, RED, REM

Duality Model of TCP/AQM

Any link algorithm that stabilizes queue

- generates Lagrange multipliers
- solves dual problem

Primal-dual algorithm:

$$x(t+1) = F(p(t), x(t))$$

Reno, Vegas

$$p(t+1) = G(p(t), x(t))$$

DropTail, RED, REM

Overview

Protocol (Reno, Vegas, RED, REM/PI...)

$$\begin{aligned}x(t+1) &= F(p(t), x(t)) \\ p(t+1) &= G(p(t), x(t))\end{aligned}$$

Equilibrium

- Performance
 - Throughput, loss, delay
- Fairness
- Utility

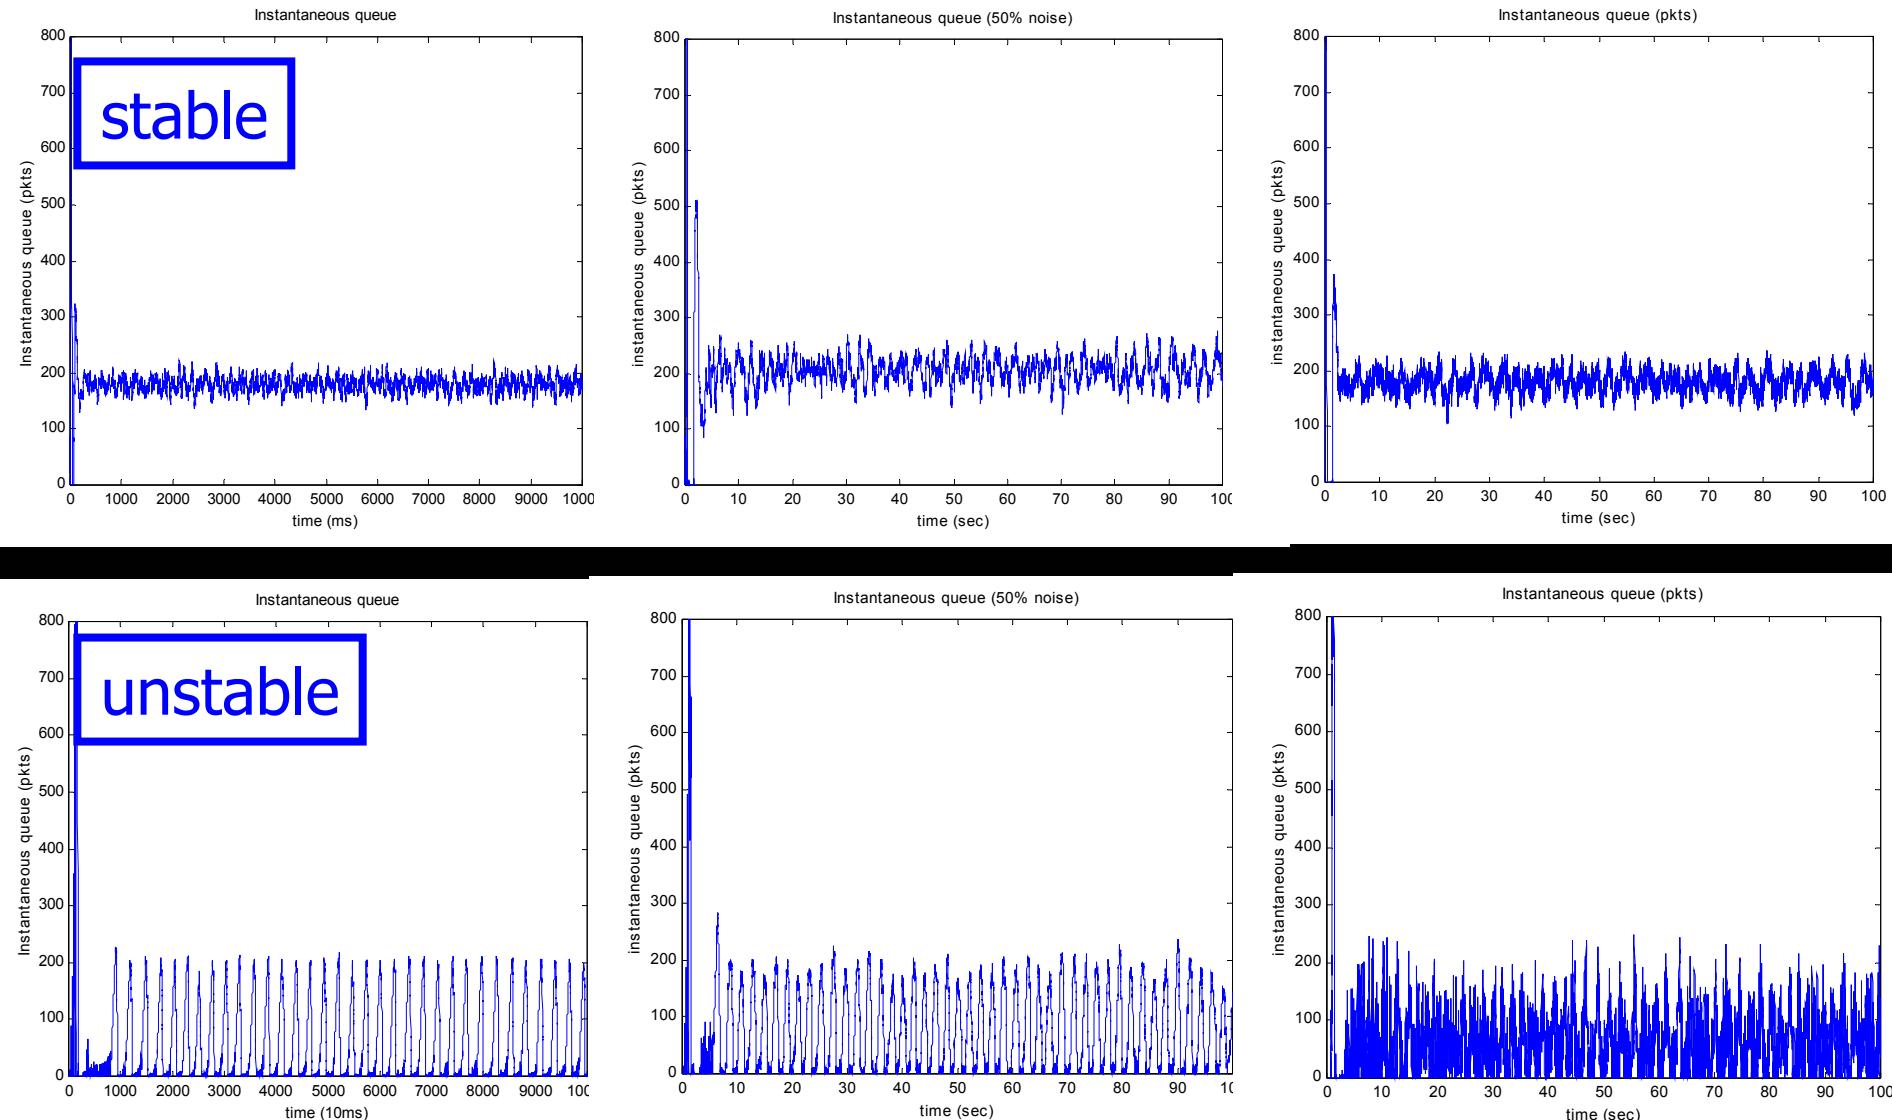
Dynamics

- Local stability
- Cost of stabilization

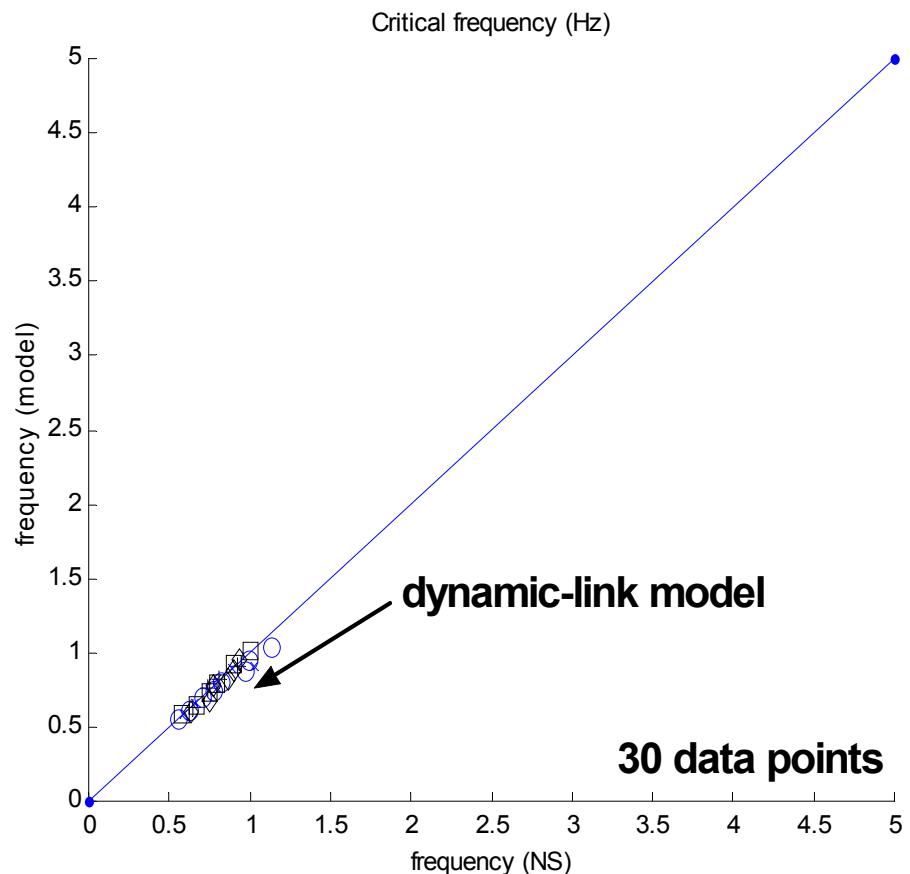
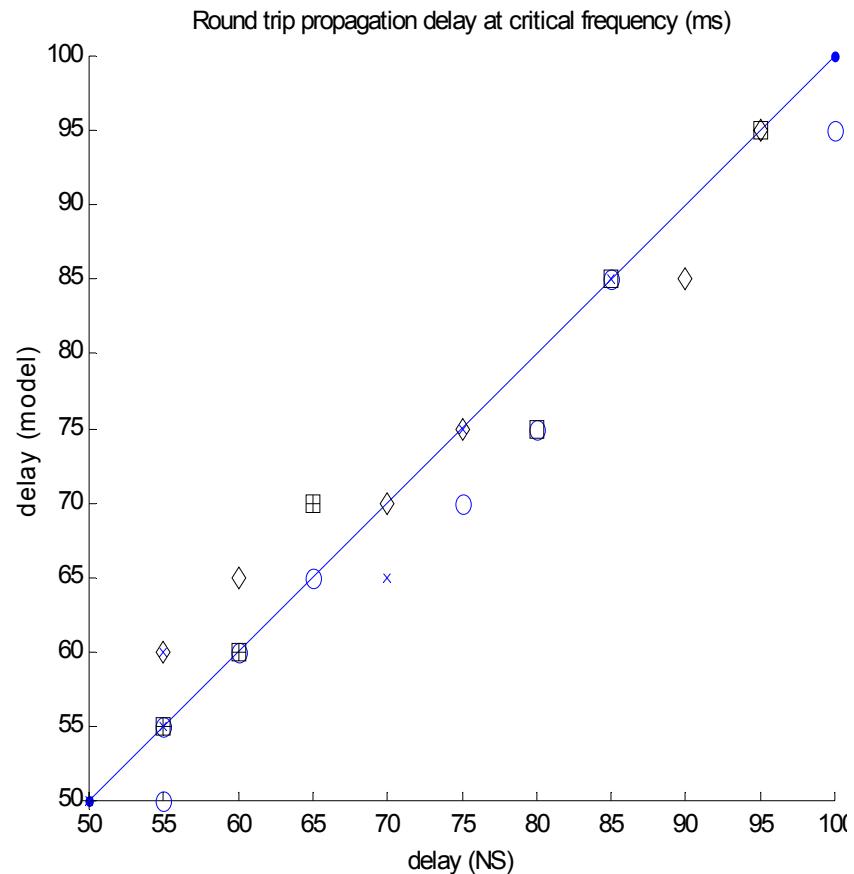
Dynamics

- Small effect on queue
 - AIMD
 - Mice traffic
 - Heterogeneity
- Big effect on queue
 - Stability!

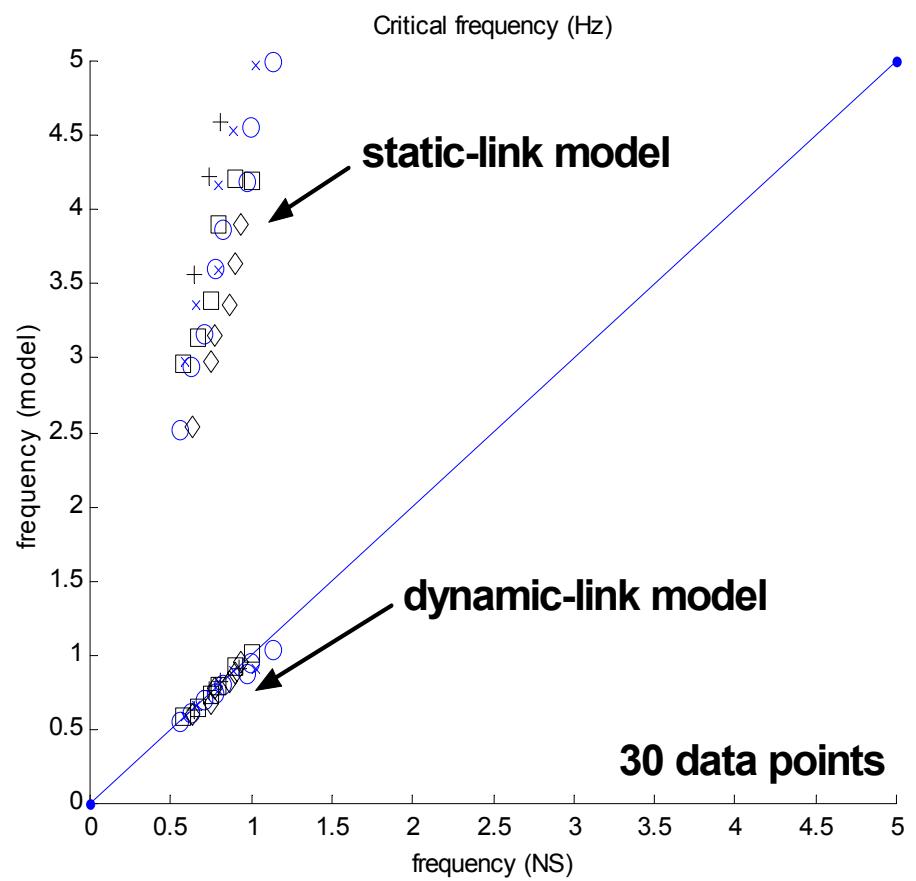
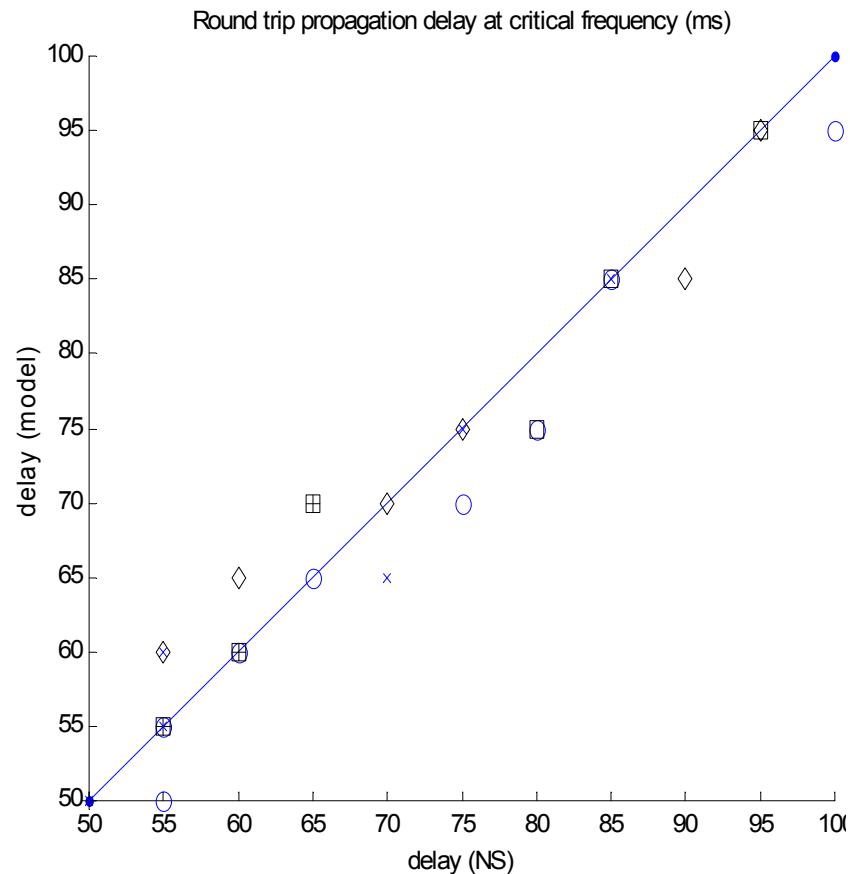
Protocol stability



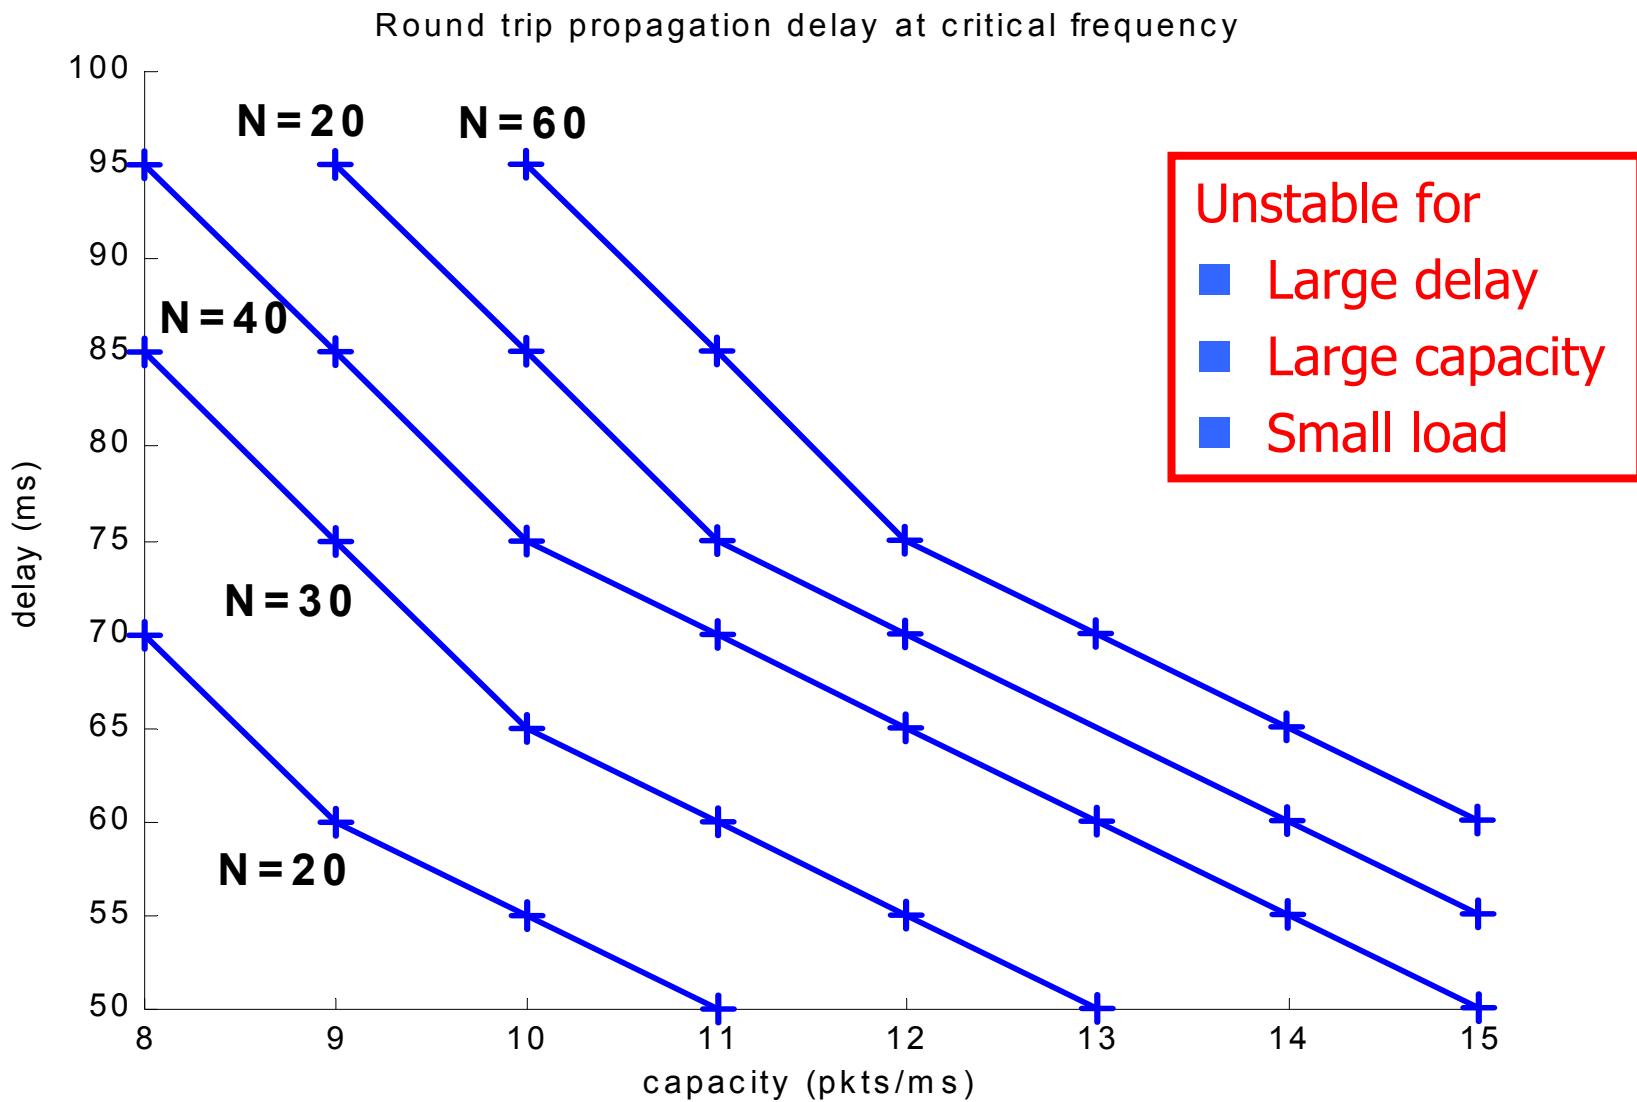
Validation



Validation



Stability region

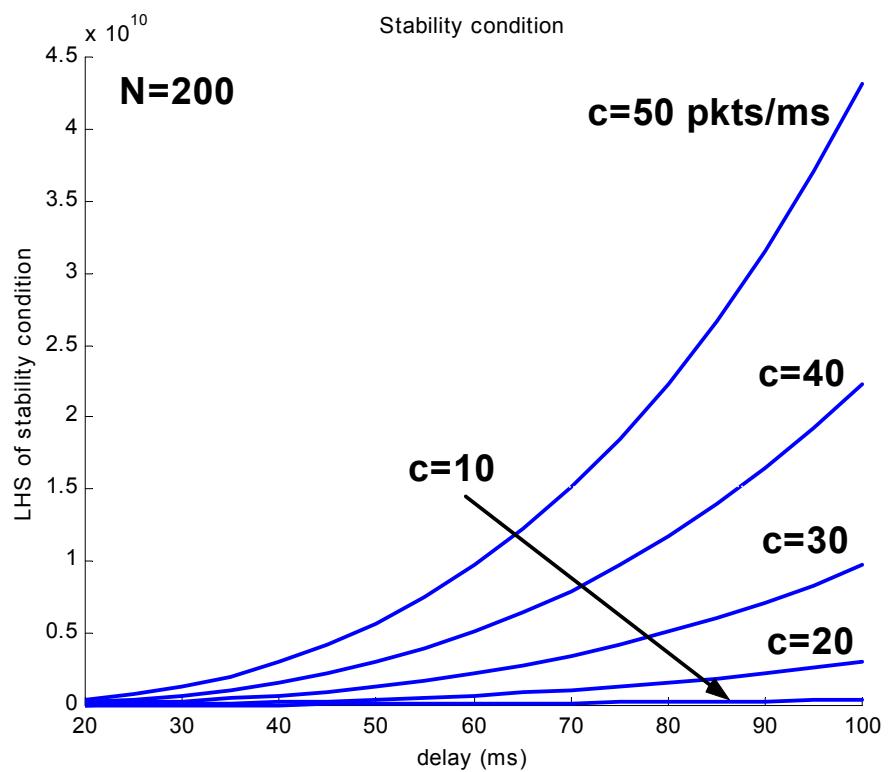
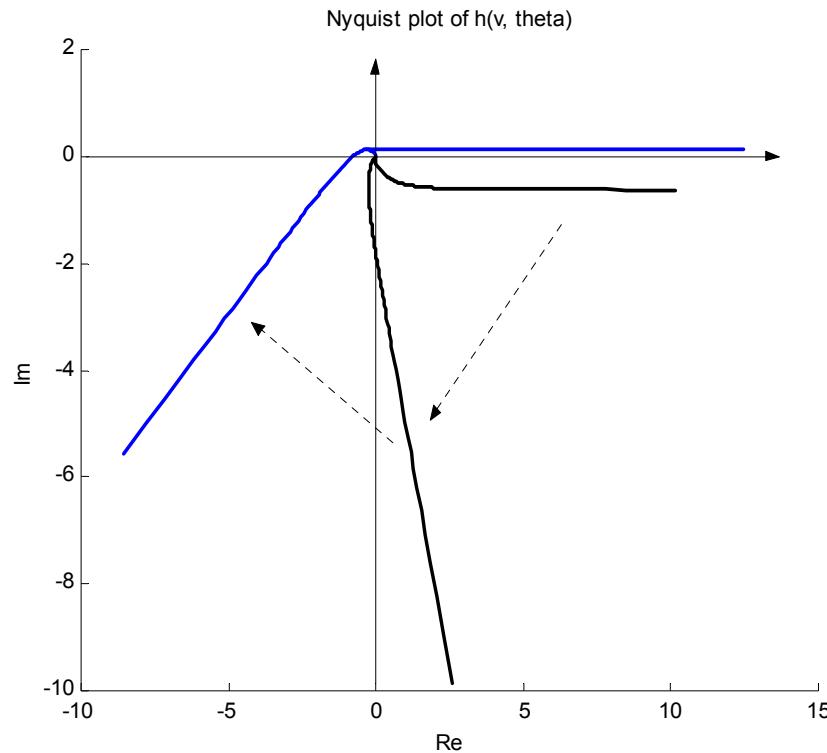


Stability condition

Theorem

TCP/RED stable if

$$c^3 \bar{\tau}^3 |H| \leq \frac{1 - \beta}{\rho}$$



Outline

- HOT traffic
- Duality model of TCP/AQM
 - Current protocols (Reno, Vegas, RED, REM/PI, AVQ...)
- FAST control
 - New protocols
- Theory of complex systems

Scalable control (Paganini, Doyle, Low 01)

TCP

$$x_i(t) = \bar{x}_i e^{-\frac{\alpha_i}{\tau_i m_i} q_i(t)}$$

AQM

$$\dot{p}_l(t) = \frac{1}{c_l} (y_l(t) - c_l)$$

Theorem (Paganini, Doyle, Low 2001)

Provided R is full rank, feedback loop is locally stable for arbitrary delay, capacity, load and topology

FAST: Fast AQM Scalable TCP

netlab.caltech.edu/FAST/

- Primal approach (Cambridge, UIUC/UPenn)

$$\text{Source: } F_r(s) = k_r \frac{1}{s\tau_r + a_r}$$
$$\text{Link : } G_l(s) = k_l \frac{s+1/T}{s}$$

- Dual approach (Caltech/UCLA)

$$\text{Source: } F_r(s) = k_r \frac{s+1/T}{s\tau_r + a_r}$$
$$\text{Link : } G_l(s) = k_l \frac{1}{s}$$

FAST: Fast AQM Scalable TCP

netlab.caltech.edu/FAST/

Benefits:

- Linear stability scalable with
 - delay
 - capacity
 - routing
 - load
- Low loss and queueing delay
- High utilization
- Arbitrary fairness

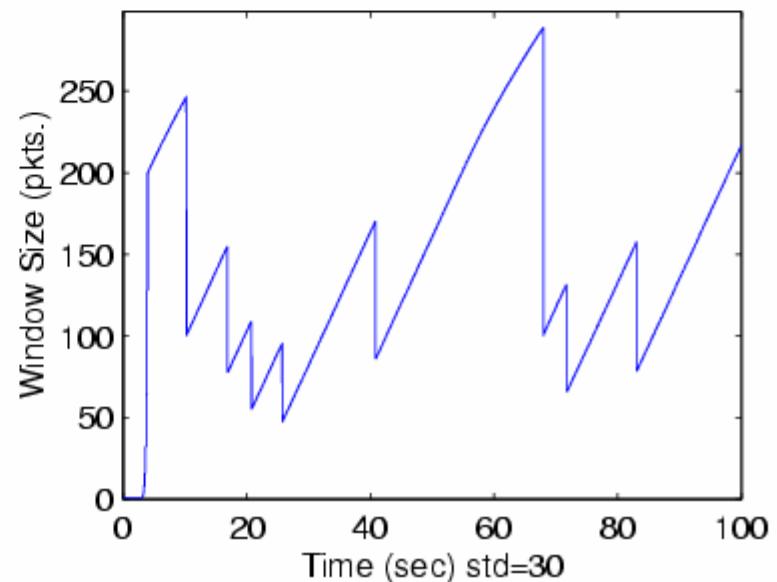
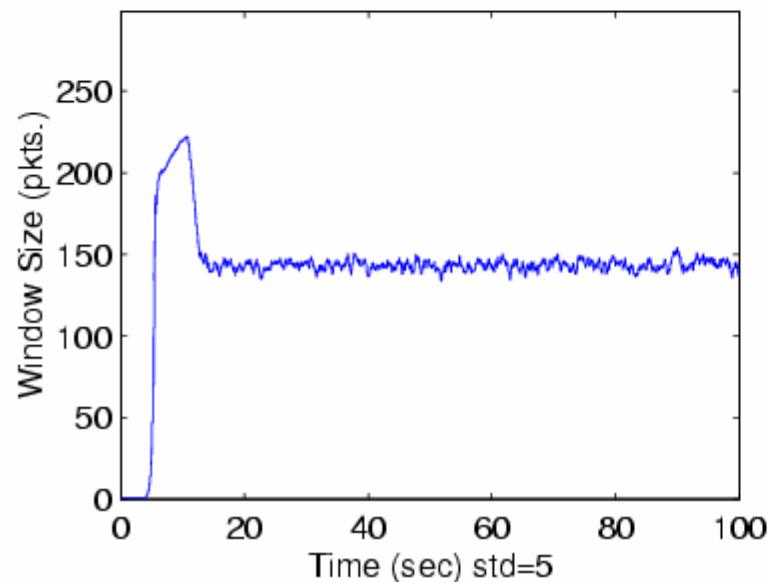
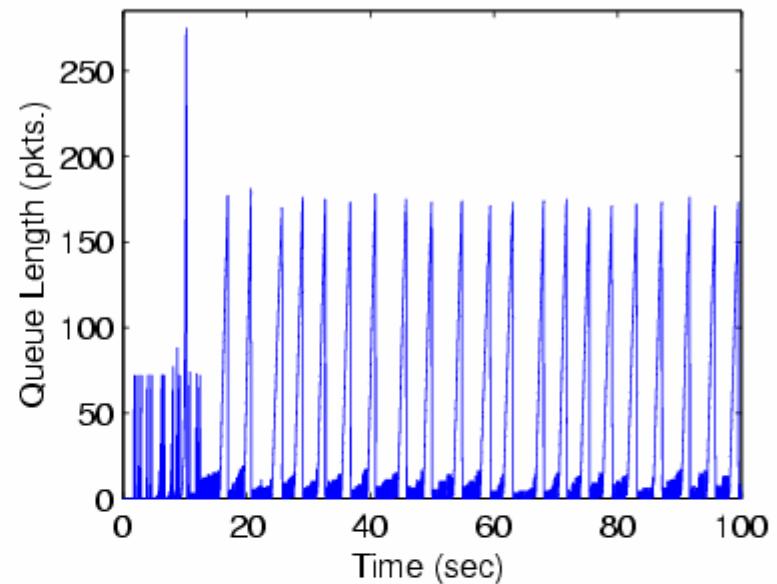
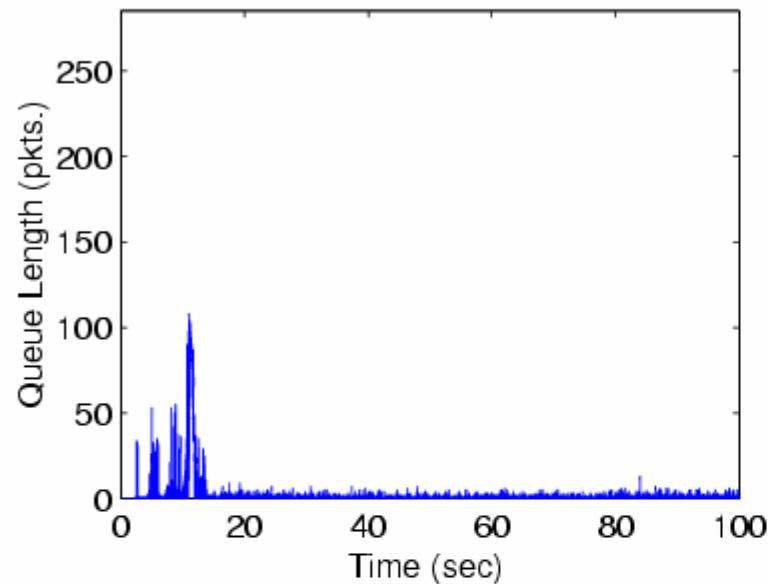
FAST: Fast AQM Scalable TCP

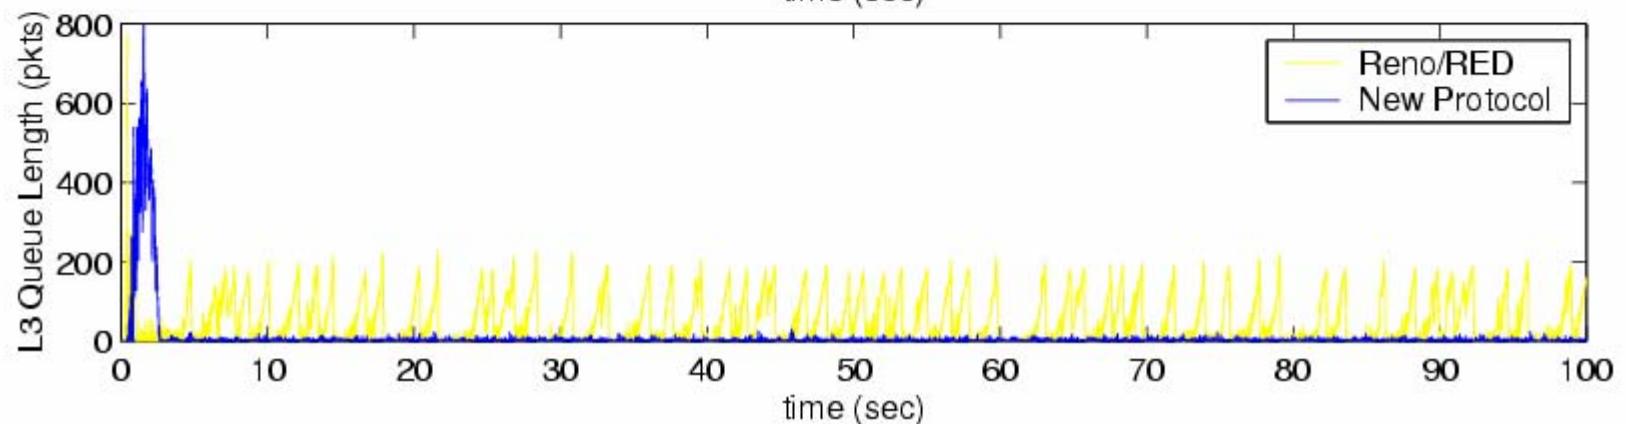
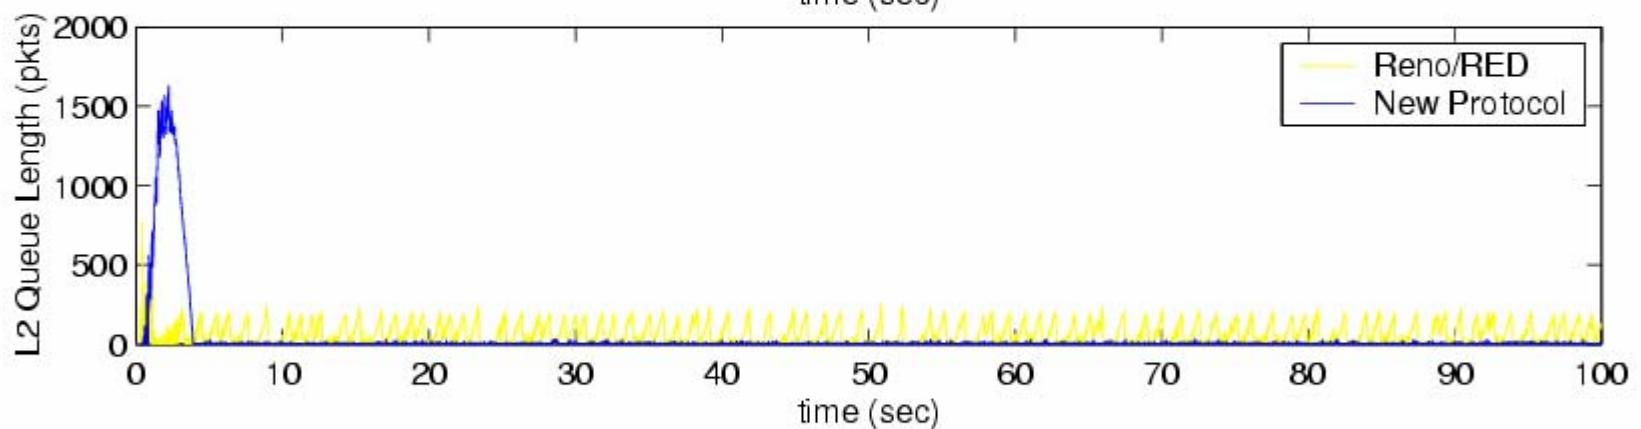
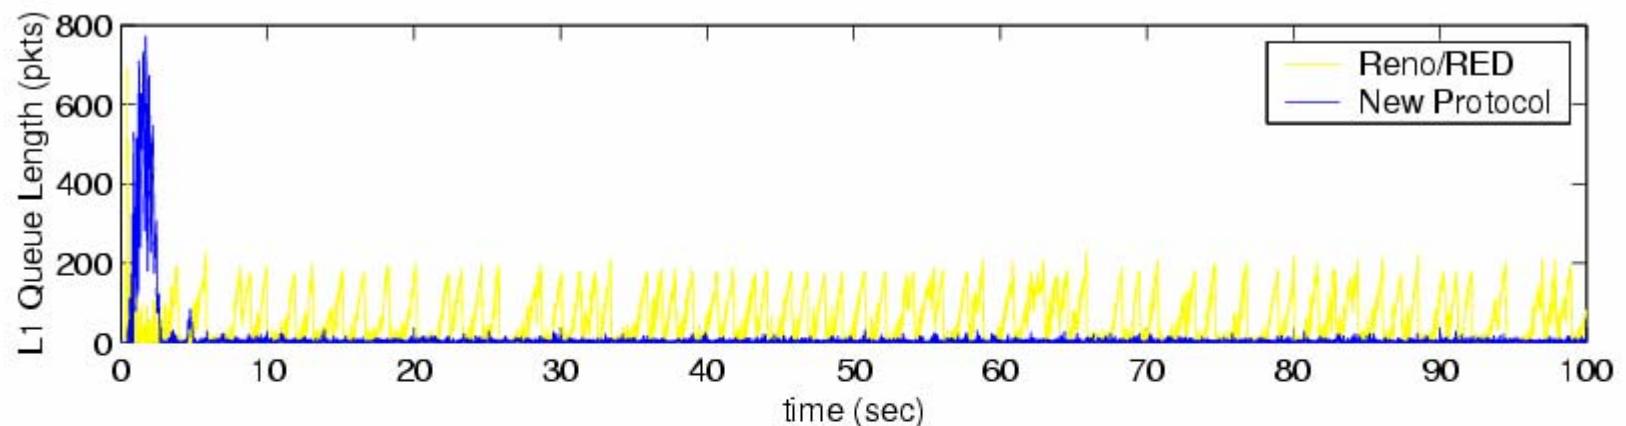
netlab.caltech.edu/FAST/

Status:

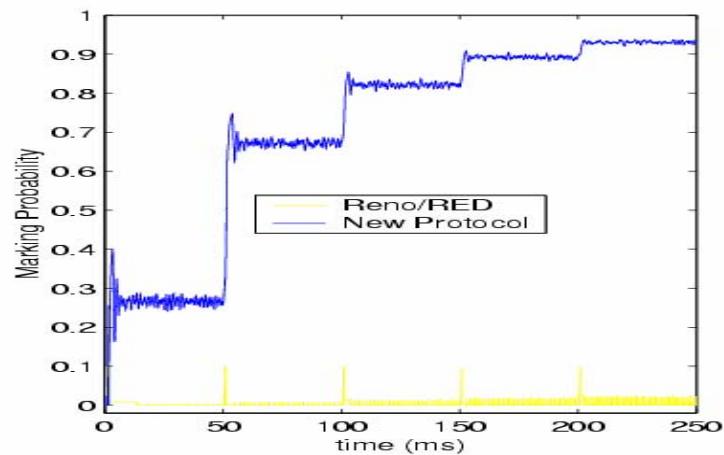
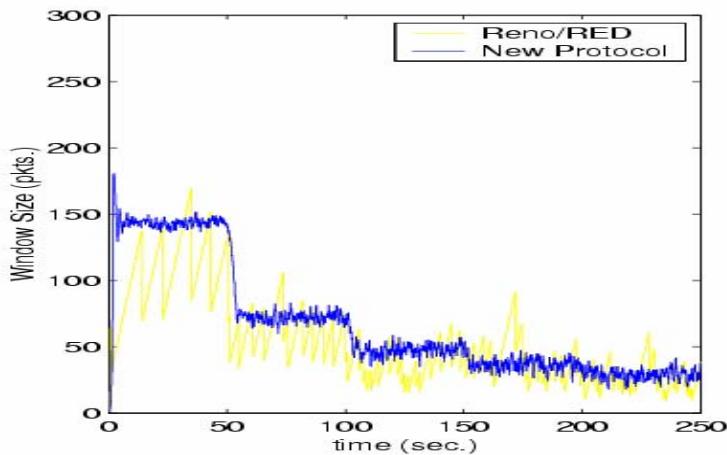
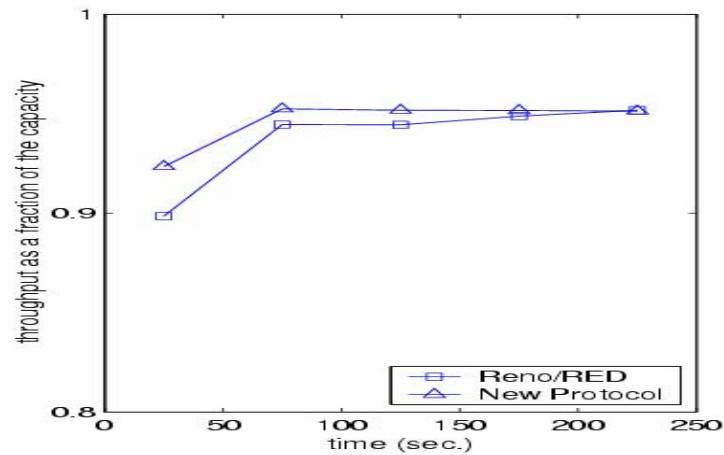
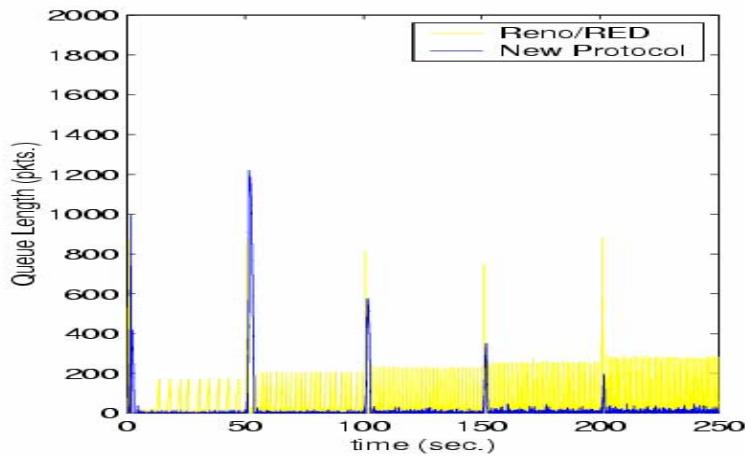
- Convergence of primal and dual approaches
- Implementation issues
 - How to generate fluid prices at links
 - How to estimate e2e prices at sources
- Quantify improvement through simulations
- Incremental deployment strategies

C=25pkts/msec, rtt=120msec, #sources=20, NewProtocol C=25pkts/msec, rtt=120msec, #sources=20, Reno/RED





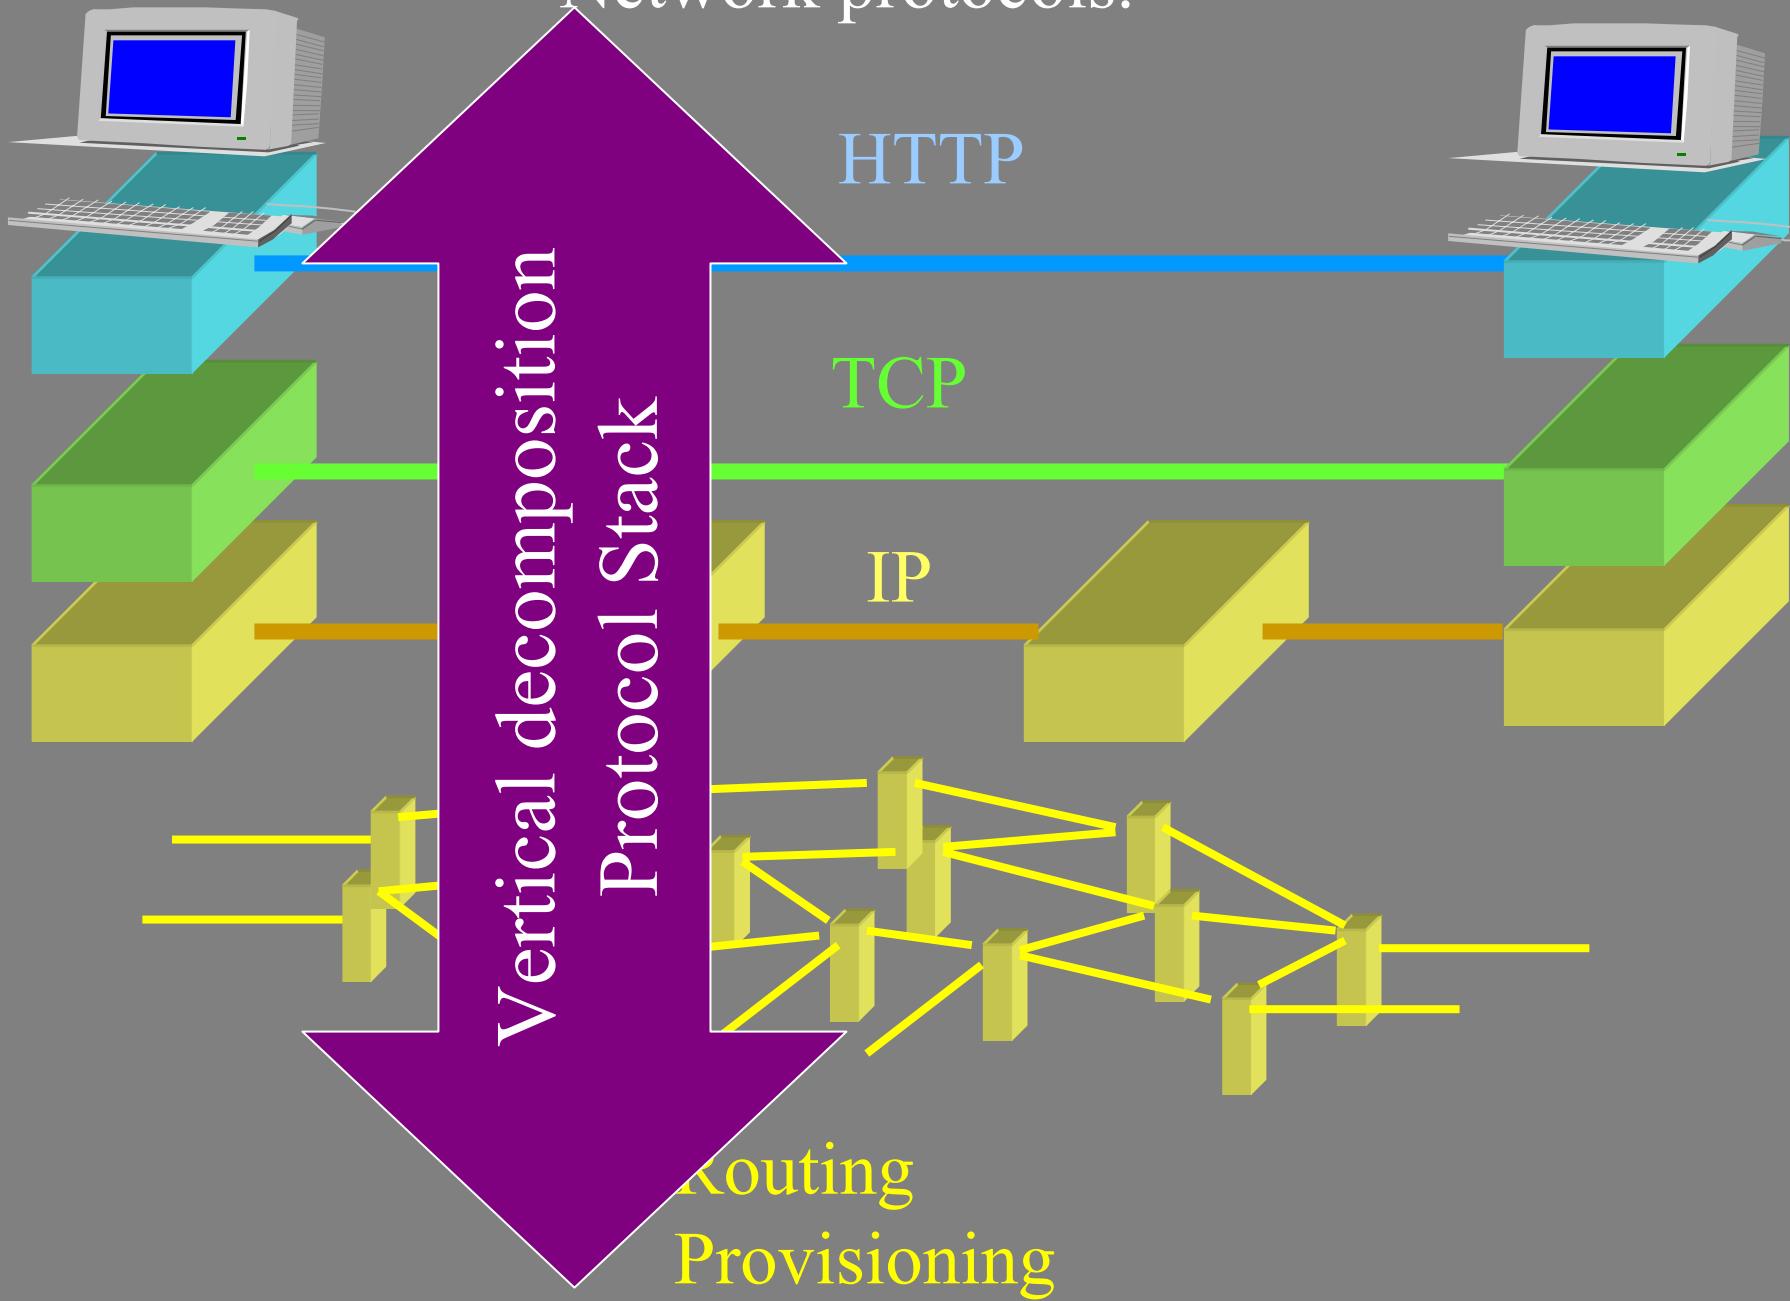
Dynamics



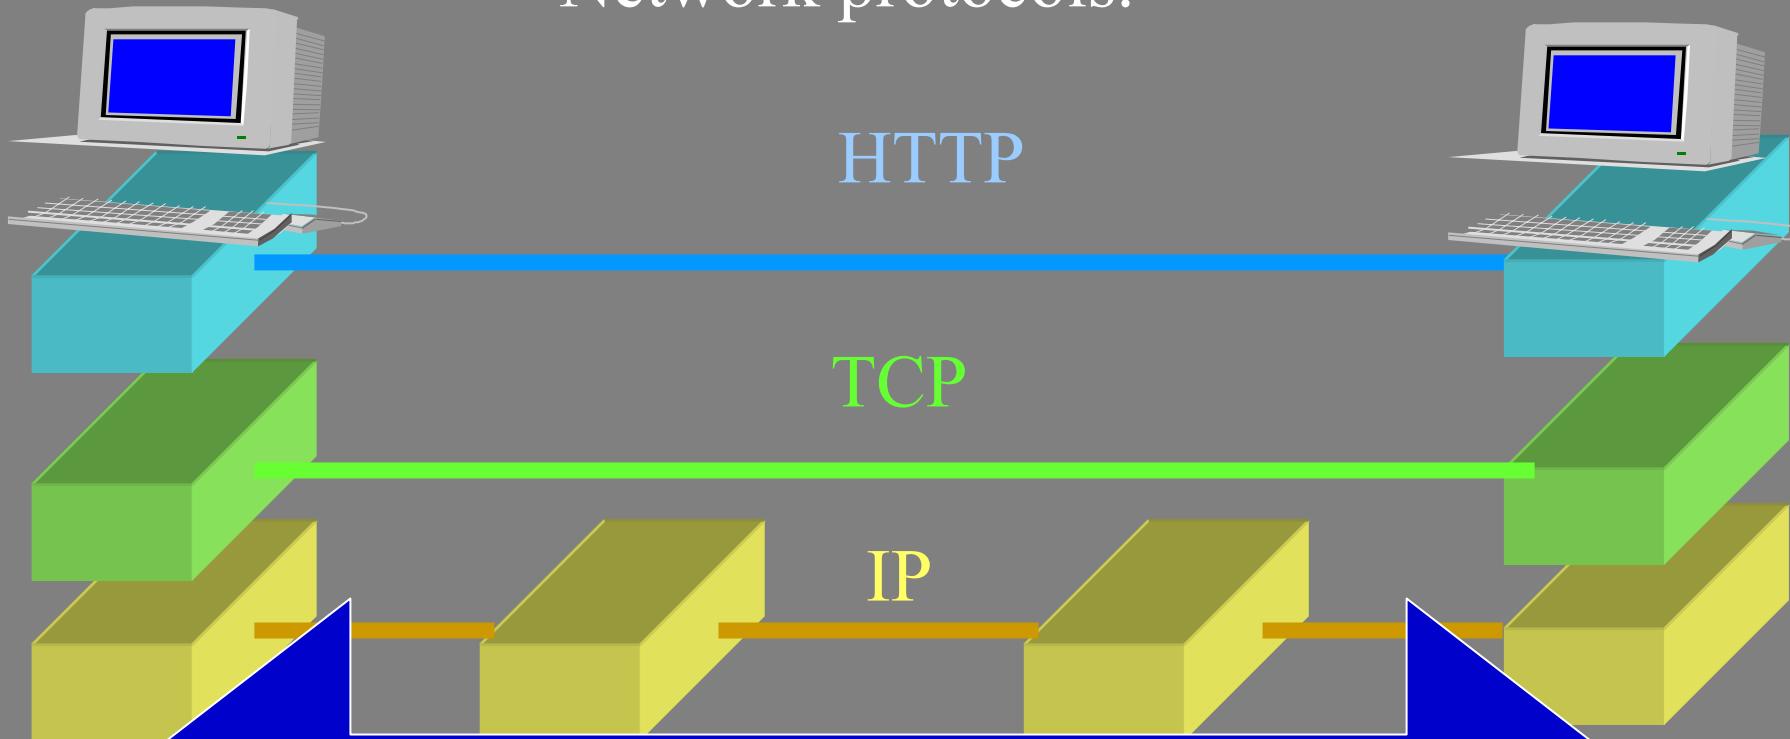
Outline

- HOT traffic
- Duality model of TCP/AQM
 - Current protocols (Reno, Vegas, RED, REM/PI, AVQ...)
- FAST control
 - New protocols
- Theory of complex systems

Network protocols.



Network protocols.



Horizontal decomposition
Each level is decentralized and asynchronous

Routing
Provisioning

Key elements of new theory

- HOT traffic
 - Source coding into mice and elephants (Doyle)
- Primal/dual vertical and horizontal decomposition
- FAST control
 - Low loss/delay, high utilization, fair, stable
- How bad is IP routing in a “well-provisioned” network?
Conjecture: Not bad.
- Vertical and horizontal integration can be made “nearly” optimal in an asymptotic sense.
- Lots of people are working out details
 - Kelly, Doyle, Low, Paganini, Srikant, Towsley, Vinnicombe, ...